Daniel Ellsberg
Abstract of: "Winning At Russian Roulette", read at September meetings of the Econometric Society

The von Neumann-Morgenstern utility axioms are norms of behavior for choice under risk, i.e., when the probabilities of possible outcomes are known. For choice under uncertainty, when probabilities are unknown or undefined as in games or "games against Nature", various alternative decision-rules have been proposed, including minimax, minimax regret and the Hurwicz criterion. In certain situations that combine both risk and uncertainty, these two approaches to "rationality" may conflict.

The utility axioms imply restrictions on "rational" preferences among risky prospects ("lottery tickets"): sets of alternative outcomes ("prizes") with known probabilities of occurrence. The situations we consider are prospects that offer as outcomes strategies in a "game against Nature." A lottery ticket with strategies for prizes (thus combining risk and uncertainty) corresponds to the notion of a mixed strategy.

One axiom requires that a lottery ticket will never be preferred to all its component outcomes. But a man who chooses among strategies (acting under uncertainty) according to the Wald minimax criterion will often prefer a mixed strategy to any of its "pure" strategy components. Minimaxing regret or the Hurwicz criterion will lead to similar violations of this and other axioms. Corresponding results follow if the outcomes are strategies in a game against a reasonable opponent.

Thus, "reasonable" men who minimax, etc., under uncertainty will, in these situations, serve as counterexamples to the utility axioms. The logical conflict here cuts both ways: it may be regarded as imposing a limitation on the utility axioms, or as evidence of the "unreasonableness" of the decision-rules in question. But it cannot be easily dismissed as applying "merely" to prospects that combine both risk and uncertainty: in reality, what prospect does not?

Consider the Russian Roulette fan. In preferring the mixed strategy (Life, Death; 5/6, 1/6) to either of its components, he violates the utility axioms. Yet, taking account of the uncertainties associated with these "outcomes," and assuming that the player minimaxes regret, such a choice may be precisely optimal for him. Indeed, the same approach suggests that certain players should rationally put two, three or four bullets in the chamber before randomizing.

1. Keystone of ec theory is theory of action, decision.

2. Ec model has been theory of the rational choice, the deduction from premisses, the weighing of alternatives. Hart, Ramsey.

3. But if we don't know "what difference it makes"? Should we modify model or abandon it?

4. Alternative approaches.

5. Why was decision approach attractive? a) appeal to intuition; what would a reasonable man do? b) the data was available, easily, for the problems we wanted to solve; ~~we wanted~~ simple aspects of consumer preference, profit--Quantitative data, no apparent problems of measurement; c) direct relevance to normative questions.

6. Confronted with "lack of realism", first impulse is to retain model as an approximation; it was always approximate, and we know that in many problems it is a good approximation. But the critics were concerned with more than the mere departure from reality of the theory in the areas where it worked best. Attempts to apply the theory, assuming that the man acted"as if" he knew the relevant data (Alchian) run up against this: a) when you depart from hypothetical data, how can you apply the theory, make predictions? Where is the data that corresponds to the data in the model? The theorist is no more certain than the actor. What are the appropriate expectations? Quote Friedman, Stigler. Two different theorists, with this approach, will make different predictions. b) We need theory of individual decisions, key choices (vs. Alchian);...and we know that there are significant differences in reaction even to given experience, with identical preferences; even in the large, many phenomena (profit?) are based directly on the fact of differing opinion and would be hard to explain on a theory that assumed men acted "as if" they had "rational expectations." c)

c) Knight: certain institutions cannot be well explained in a world of certainty--information-gathering, research, organization; or the differing expectations that lead to profit. Moreover, certain decisions do not seem to be explainable to definite expectations at all; return to this.

7. Thus, standard theory runs up against ~~problems~~ difficulties of application, and seems quite unable to cope with certain problems. Still, all the approaches I will consider do retain the framework of the decision, with modifications. There are reasons: 1) we still want theory of individual decisions, and we want to predict differences in reaction; 2) The approach still seems appropriate to economic and organization problems, and it does enable us to use our introspection and casual observation at least in the judging and framing of hypotheses for testing We will not devote our first efforts to hypotheses that seem to be contradicted by such casual testimony.

8. However: a) we will have to give up hope that preferences, even of the businessmen, can be inferred directly by introspection, or that opinions, expectations can be inferred simply and directly from data on observable economic magnitudes. We will still look for simple hypotheses, but we need not expect them to come quickly and easily. Schumpeter's view. ~~But~~ Mechanical inference from time series does not account for differences (in time and between persons). Knight, Shackle on expectation. The moral is: look first for good relationships between expectations and decisions, not between experience and expectations; i.e., take expectations as data. This is not a counsel of despiar, because the possibility of observing expectations directly, and prefs, is much greater than it used to be. But it does get us into measurement problems. The predictive power to which we can aspire will de-

on the extent of our knowledge about the data. For example, with only ordinal knowledge of preferences and likelihoods, only very simple hypotheses are possible: e.g., ~~that~~ that the man acts as though the most likely ~~hy~~ outcome were certain, or that the most pleasant, or worst outcome is certain. But most agree that useful hypotheses need stronger restrictions on data; Suppes and Davidson. But how to measure degree?

8. In the ec literature, I will consider Knight, Shackle, Ramsey-Savage, Wald-Hurwicz-Savage, and von Neumann. Each of them works in the general area staked out by Knight (exception, von Neumann), and each of them claims the whole of it for his model--at least, for any pair there are certain problems that both claim; and they generally declare their models to be incompatible. Therefore, before committing ourselves to full-scale application of any one, some thought is called for in analyzing and comparing them. What sorts of experiments and evidencewould enable us to distinguish between the applicability of the two models--i.e., in what meaningful way are they incompatible; what evidence would be useful in several approaches--to what extent are their empirical implications the same? Would further analysis of the problem be useful; are there reasons to suspect that one approach is more suitable in certain particular areas and another in others? ~~Are there~~ Is there evidence of a quick and easy sort suggesting that one approach is more suitable than another in particular problems?

9. It is this last sort of "plausibility" that I will invoke, but only tentatively. There has been great emphasis on plausibility of various hypotheses, their deducibility from "obvious" axioms. Partly this is a carryover from a period when scientific technique and sources of data were primitive, when it was very hard to decide even

I propose to explore the hypothesis that in certain, specified circumstances, degrees of belief, insofar as they can be inferred from choice, will be ~~interdependent~~ dependent on payoffs; that where this influence is marked, models that are formally non-probabilistic are the most convenient approximation; and that it is precisely in these cases that the non-probabilistic models are ~~most~~ at their best.

Thus, I offer a theory that relates the two types of model to each other, that prescribes conditions under which decision-making behavior undergoes a certain change, so that one abstract model gradually becomes more appropriate than the other, each model being each stage a more or less good approximation of the other.

This approach will be valuable if it leads to more fruitful experimentation, by indicating separate areas of behavior where the various models are most appropriate. But I will defend it as more than an heuristic device, a "point of view" that "suggests" certain lines of investigation that may, more or less accidentally, turn out to be fruitful. This theory itself has empirical content, with testable implications that ~~go beyond~~ lie outstide those of any present proposed model. These implications concern an as-yet-unexplored area of decisions that lie, as the theory implies, on the borderline of the regions where the two familiar models apply. Although empirical results may be more difficult to obtain and less definite than in the two "extremities," the successful application of the theory there would indicate that the underlying analysis, which specifies the areas of application of the two other models, is valid.

THE THEORY IS:

1. In general, the decision problem is to be analyzed in terms of acts, of expectations as to the outcomes, and of preferences over the outcomes; and of "confidence."
2. In general, expectations are not single-valued, and a measure of degree of likelihood is relevant: which may range from a ratio scale to an ordered metric or an ordinal scale; moreover, preference is assumed measurable either to a linear transformation or an ordered metric scale.
3. Both expectations and preferences are to be measured either by psychological interrogation (including Shackle's method) or by observations of choices in uncertainty-situations. At best, empirical investigation will probably determine the variables up to an ordered metric scale, but it may be possible to approximate results by a numerical interval scale.
4. Confidence is not regarded as even approximately numerical, but is ordered and has some ordered metric properties. When it is high, it will be possible to infer from choices a (roughly) numerical scale of utility and an independent numerical scale of likelihood, with choice determined by the mathematical expectation (Savage, Ramsey); these scales will correspond closely to scales determined by psychological methods.
5. ~~Confidence~~ When "confidence" is high, it is possible to infer from choices independent measures of likelihood and utility, such that

A

A: There are now a number of competing formalizations of the problem of choice under uncertainty. ~~On the whole their areas o they claim~~

They claim to cover many of the same problems; yet their proponents generally do not regard them as alternative, or claim that their system is merely more convenient, more fruitful, or broader in scope. On the contrary, it is argued that the frameworks are logically inconsistent, ~~and~~ that certain phenomena can be described in terms of certain concepts and no others, that ~~certai~~ competing concepts have entirely inappropriate formal properties to apply to the data at hand.

B. In many cases, a sweeping rejection of a conceptual approach is justified only by ~~a~~ statements of the intuitive "feel" of the concepts, a quasi-aesthetic reaction; in other cases, a model is rejected because of shortcomings in a particular interpretation, though the model might lend itself to interpretations that would avoid these objections. ~~Actually, in most of these cases models, the~~

C. In general, the empirical implications of the models (with given interpretations) have rarely been spelled out in a fashion that allows us to discriminate between ~~them~~ *models*. (though we can readily reject particular interpretations of a model on the basis of experience). What evidence would lead us to conclude that the subjects did *not* act *as if* they assigned numerical probabilities to events? What observations would *contradict* the hypothesis that degrees of belief ~~acted like numerical probabilities~~ were additive? Not even the critics of the various proposals have produced such tests.

D. Neither the critics nor the defenders of the various schemes have made much attempt to delimit the area of application of a given approach. Each approach is able to exhibit concrete examples where its approach is suggestive or validated; but in each case there are other concrete examples where it does not seem appropriate, and some other approach seems better. ~~These unfavorable examples, however, are used by each approach to di reject each other approach in toto~~ There is a marked division of labor, however, inthe analysis of these counterexamples. It is very rare, except in footnotes, to find the proponent of one scheme discussing examples that are unfavorable to his own approach, or the strengths of some other approach. Such a discussion might well lead to an attempt on his part to analyze specifi conditions under which one or another approach is to be preferred. The preoccupation by each writer with the strengths of his approach and the weaknesses of all others has discouraged such an analysis.

E. The attempt to derive various hypotheses from "axioms" does not affect this conclusion. Since their main purpose is to lend plausibility to the "theorem," the individual axioms are chosen so ~~as to~~ that ~~each one~~, so far as possible, each one appears unlikely to be violated. This applies particularly to their statement, which is almost never in a form which will suggest immediately conditions under which violations might ~~be expected~~ occur. ~~And yet, no author denies that violations will occur~~ Examples of violations are chosen and discussed in a way that makes them seem, somehow, not serious, not worth considering. Yet every author readily admits that the theorem, the hypothesis, will be verrified, at best, only approximately, and may not be verified at all. Every empirical attempt so far to apply the hypotheses has confirmed this. ~~Yet~~ This means that some "axioms" are being violated. Which ones? When? How serious are the violations, relatively? The discussion of the axioms gives no hint as to how we might expect these questions to be answered. Moreover, it is significant that none of the empirical investigations has attempted ~~to~~ even to categorize violations in terms of the axioms, let alone to investigate conditions under which various violations occur. This suggests that these axiom systems make no contribution to the actual empirical investigation of the hypothesis (though, by strengthening its "plausibility" they may have led to its being considered worthy of investigation).

Moreover, the atmosphere in which these hypotheses are discussed has led to tests which compared one approach overall to another

2.80
6
16.40

the subject maximizes the math. expect of utility. When confidence is low, relative to the scale of payoffs, the likelihoods will be interdependent with payoffs; this can be seen by ~~raising~~ shifting the payoffs between events, or raising the level of payoffs (the utilities having been previously determined); or by take events with the same likelihoods (as indicated by experiments with small payoffs) but greater confidence, and the same level of payoffs; or by comparing with results of psychological measurement of likelihoods--close correspondence expected when choices involve small payoffs or high confidence, but deviation when choices involve large payoffs and-or low confidence.

EATON'S BERKSHIRE BOND U S A BERKSHIRE

he does imply ordinal likelihoods

In what form are expectations? Knight doesn't say./ Why not consider them numerical probs? Partly because that seems too precise, mainly because K thinks of numerical probs as derived only from well-defined randomized experiments. Indeed, the main reason for imagining that degrees of belief might act like probs is that where frequencies are known, we expect degrees of belief to be closely related to them; and Knight does expect this--on an individual basis; moreover, he expects them to operate much like probs, in cases of grouping. Quote Arrow.

But a result of lack of knowledge, arbitrary grouping is lack of confidence. Can this confidence judgment be expressed in probs? No, says Knight. Keep this in mind as a possible phenomenon which weighs against numerical probs. He clearly can't conceive of numerical probs which aren't derived from frequencies, or " a priori" calculations. How measure them? Where do they come from. (they might be his expect)

Shackle uses same analysis, ~~again,~~ more than the availability of frequencies, he stresses their relevance to the unique, or crucial or isolated trial. This leads him to reject numerical probs--again, he can't really conceive of numerical probs being meaningful apart from frequencies--though again, he himself proposes operations that could make such probs meaningful.

Shackle does provide a formal structure for the decision and offers an operation for measuring expectations. p.s.

	Confidence	
Payoffs	Low	High
[illegible]	"	numerical probs + utils
High	Non-prob models	Expect. dependent on payoffs

There's an upper limit to confidence

Hence difficult to get utils for high payoffs except for very high confidence using whether it worked well or not, and when the main application was on hypothetical data. Nowadays, I feel it is inappropriate to spend the same effort in deciding whether or not a hypothesis is plausible, reasonable, convincing: we should try it and find out how it works. Some effort is worthwhile for purposes of framing useful hypotheses and economizing our resources in testing them; but this is only the first step.

Another reason for such concern with the "reasonableness" of the hypotheses is due to a desire to arrive at assertions that will still simultaneously describe and prescribe. There really seems no reason to insist on a theory that does both well. And to those who are interested mainly in the normative aspects, I suggest that their preoccupation is premature. Economic advice must rest on knowledge, and knowledge of actual decision-making. In the absence of sound ability to describe and predict how people actually do make their choices, our advice will remain without substance, purely formal.

Therefore, I propose to look on the current discussion of rational decision-making as a source of testable hypotheses about actual decision-making--possibly over the protests of some writers.

10. Knight made a basic distinction between problems where knowledge of frequency-ratios existed and was relevant. Relevance required that the ~~problems~~ decision concern the total results of a group of trials; but K. didn't stress this; where ratios were known generally, insurance was possible or self-insurance. The main problem was lack of knowledge of frequencies; problem of valid grouping of instances. Quote. The results of this were: a) expectations would differ; b) hence, insurance would be impossible;and individual would act under uncertainty facing the single act; c) lack of confidence--neither single-valued expectations nor , in some sense, confidence.

Uncertainty

1. The keystone of economic theory is a theory of economic decision-making. Those economists who aspire to an "astronomy of commodities" cannot evade the fact that economic forces work through individuals, that factors which influence economic variables do so by influencing human action. For certain purposes, we may take certain key decisions as given, as autonomous for the purposes of the problem; but unless we are able to predict a mass of ~~predictions~~ resulting decisions, we will be able to say very little about the effects of these key decisions, or of anything else.

2. There are many possible approaches to the problem of predicting human action. It can be regarded as random, within certain limits and subject to various statistical assumptions about distribution, auto-correlation,etc; as imitative, represented by models of contagion; as based on tradition or authority, the theorists' problem being to locate the leaders or the ancient laws; as impulsive, the impulse being triggered by esoteric cues in the environment; as a learning process, partly experimental, partly influenced in a mechanical way by recent past experience.

3. Instead, economists have chosen in the past to analyze action in the economic sphere in the framework of the rational decision; their model of action is the reflective, reasoned choice from alternatives actions in the light of their consequences. Obviously, economists are aware that this does not describe "realistically" the conditions for most economic choices; but for a first approximation, they had ~~good~~ reasons for regarding an economic decision as, as Herbert Simon has put it, a "deduction from premisses."

a) First, with relevant data very scanty and hard to procure, and limited techniques or opportunities for experimentation, there was a great premium on a conceptual scheme which could make full use of introspective evidence and casual observation. It is difficult to determine or introspection ~~or conversation~~ the influence on~~e~~ one's actions ~~th~~ of tradition, authority, unconscious impulse, imitation, or ~~random~~ stochastic elements, but it is relatively easy to analyze the process of conscious deliberate, choice ~~decision~~: to categorize the "premisses" ~~elements that influence the outcome,~~ on which the choice is based, to "explain" the conclusion. On the basis of reflection and everyday experience the economist was convinced that he could specify the data needed and the relationships to predict deliberate choice.

b) But ~~why expect this to be an adequate~~ on what grounds could one ~~discount~~ discount or ignore entirely the influence of the other factors we have mentioned? Notice that when this approach was new

This approach leads us to emphasize the significance of two types of data, which have been well characterized by Frank Ramsey. In espousing ~~this~~ the approach, in "Truth and Probability," he remarks:

> "I propose to take as abasis a general psychological theory, which is now universally discarded, but nevertheless comes, I think, fairly close to the truth in the sort of cases with which we are most concerned. I mean the theory that we act in the way we think most likely to realise the objects of our desires, so that a person's actions are completely determined by his desires and opinions."

Desires and opinions: these, then, are the premisses from the rational action is deduced, the data on which we propose to base our predictions. He also suggests the second reason why the economist is attached to this approach in preference to others: "it comes fairly close to the truth in the sort of cases with which we are most concerned." Specifically, in many economic decisions the role of conscious deliberation of desires and opinions seems relatively greater than in other spheres. More important, we are concerned with decisions made in and by *organiaations*. Organization behavior seems to fit the model of rational decision-making peculiarly well. Goals of the organiation tend to be integrated and explicit; estimates are in writing and the decision-making procedure is formalized; there is deliberate effort to eliminate influences or limitations that inhibit goal-achievement; and perhaps most important, the organization gathers and distributes the information relevant to action and coordinates the expectations of its members.

Finally, the same model satisfied the demands made upon economists for *advice* as to economic decisions. When someone wanted to know the "best" thing to do in a given situation, the economist had the answer right at hand; it happened to be the same assertion he was using to *predict* empirical behavior. In fact, the normative and the descriptive aspects of ~~thexxxeconomicxas~~ the assertions of economic theory have been interrelated: to the extent that they were reliable empirical generalizations, it was because people *tried*, and not merely tended, to obey rational principles.

1. The keystone of economic theorizing has always been a theory of decision-making. Not merely a theory of action; a theory of decision, quasi-conscious, emphasis on choice among possibilities; choice as a deduction from premises (Simon).

2. Why did theory develop this way? a) theory of action is essential, because most ec variables depend on at least one link of human action b) because of a lack of empirical data and even empirical methods in providing a theory of action; an approach to action as choice means we could use introspection and casual observation of mental processes of our friends as a source of hypotheses and data; c) the predictions such an approach provided as approximatations to reality at the same time doubled as advice to those who accepted rational norms; two aspects were interrelated--to the extent that conclusions were reliable, it was because people tried, not merely tended, to follow them.

3. Even where it did work descriptively, ~~this was~~ the predictions did not follow simply from assumption of maximizing behavior--which in it-sel has very weak or non-existent empirical implications. It worked where our introspective knowledge of preferences and expectations was most detailed and reliable; in particular, business decisions. Profit; repetitive decisions, organization behavior, information gathering, physical productivity laws.

- Doesn't apply under uncertainty.

4. Econos have stressed one aspect or the other depending on their tastes for advice or prediction. But when it doesn't work, issue is raised: should it be scrapped or modified? Or should we retain it for normative purposes only? We must ask: Why doesn't it work? a) It may be because actors lack the data the theory assumes, so that they are unable to follow it; we must recognize that the particular actor who comes to us for advice will probably be under the same handicap; he will be unable to use our advice for the same reasons that other fail to conform to our predictions; b) even if we have information others don't, or even if they have information but aren't using it--economic advice must be based on knowledge, predictive power, and this requires a good theory of actual decision-making; i.e., if our theory can't predict actual decisions, then it isn't much good for advice either, because good advice presupposes the ability to predict other people's decisions. So we can neither predict nor advise unless we have a good theory of actual decision-making.

5. The decision approach supposes that (Ramsey) "we act in the way we think most likely to realize the objects of our desires, so that our actions are determined by our desires and opinions." This directs us to base our predictions (advice) on two classes of data, preferences and expectations (as well as telling how to utilize the data). This was convenient, because we thought we knew the preferences (at least in the crude form to which we aspired) by introspection, and the expectations by introspection plus observable data.

6. But some decisions by reasonable people could not be explained within model based on certainty (e.g., diversification) and many others could be explained better by model that assumed uncertainty. But uncertainty,not only makes old model unapplicable; it makes even a modified form of model less convenient. a) expectations and preferences can be put in a new form, but can no longer be estimated on the basis of the old data; introspection does not give plausible hypotheses on preferences among gambles, nor on expectations based on given observable data. The essential fact under uncertainty is that tastes (among

gambles, even if not among outcomes) will differ; and so will opinions. "There's no arguing over tastes" "It's difference of opinion that makes horse races" (and profits--Knight).

Naturally, many econos resist an approach that threatens to make their old source of data inapplicable. Scumpeter, on expectations. Or simply throw decision approach out: Keynes; emphasis influence of conditions on actions, hoping that individual differences will cancel out or will be dwarfed by massive influence of factors which are directly observable; simply use regression, correlation methods.
c) Moreover, question of how data is utilized to arrive at decision.
7. However, I will be conservative. The approaches I will describe retain the framework of the decision problem. Reasons:
a) approach still provides us with testable, plausible hypotheses, even though it no longer makes any one uniquely plausible. The question, "how might a reasonable man, or an average man, decide this choice?" still allows us to use our insight into average mental processes for the purpose at least of framing hypotheses and judging plausibility. This is an important function, even though it must now be recognized as only the first step toward empirical results.
b) We are interested in theory of the particular decision, because of the great influence of certain important single decisions (by big firms, government; and the framework of the rational decision seems well-designed to give it to us, because of the propensity to rationality of the orgnization (Simon).
c) The data the modified theory requires seems ~~les~~ more accessible now than twenty years ago; more data available, and improved psychological-sociological techniques for gathering, interpreting data on preferences and expectations. To the extent that we can "tap" expectations and preferences directly, we don't need to worry about how they are formed, at this stage (though that may come later).

8. What materials, then, does economic theory up to now provide us with in furnishing a theory of action under uncertainty? I will consider the recent work on "rational" axiom-systems as an attempt to arrive at useful hypotheses (even though authors may insist that they are interested in normative aspects).

9. First problem: how to describe and measure the basic data, the desires and opinions. Second, hypotheses on the relation of data to decision: these take the form of "decision rules."

((above: under certain circumstances--if even direct knowledge of his desires and opinions did not allow us to predict his decisions; or if desires and opinions changed so fast, perhaps because interdependent with actions, or if too inaccessible, sothat other data which could be regarded as determining opinions and desires would be a more efficient predictor; or, even if instruments chosen could be predicted, if they were quite unreliable, so that chief problem was to predict the effect of instrumental actions rather than their choice--then decision approach would not be useful. Or if choice of action is simply not made with reference to desires and opinions, but w.r.t. tradition, authority, imitation, reaction: which might or might not be reflected in conscious or ascertainable desires and opinions, but which in any case would be both causally precedent and more accessible.))

10. It is agreed that more data on preferences is necessary under uncertainty than mere knowledge of ordinal ranking; some measure of degree of preference is relevant. We will not be mainly concerned with this problem except to mention that data on strength of preference is available from two sources: a) insofar as it is a subjective magnitude,

it is accessible by psych methods developed in last thirty years for measuring all psych magnitudes; b) insofar as it is conceived as an influence on choice under uncertainty, it can be measured by its effects on that choice (thus, choices in certain situations can be used to provide data on preferences in order to predict choice in other situations.

11. Our problem: the measurement of the likelihood. Again, this is susceptibe to two approaches: a) psychological measurement; b) as it influences choice.

12. First, where to draw the line between certainty and uncertainty; i.e., where to circumscribe the application of the familiar theory of rational choice. Knight was one of the first to draw the line, i.e., to suggest that there was a large and significant class of behavior that lay outside a particular boundary. In terms of actual situations, his split has remained valid; we won't quarrel with it.

But his analysis of the conceptual basis of the split is dubious; moreover, he may have overlooked important subclasses within the area he labelled uncertainty. He defined the crucial difference as being the existence and relevance of "measurable" (i.e. numerical) probabilities. Actually, it may be possible to define meaningful numerical probs not only for the "risk" situations he described, but for neighboring areas both in his "certainty" and his "uncertainty"; and there may be a significant split within his "uncertainty" (where ~~it~~ in one class, it is not useful to define numerical probs independent of outcomes).

13. Kis risk-situations were a class of decision problems where: uncertainty about the results of a single action was unimportant, since the relevant outcomes were the total results of a long series of trials; and knowledge of such total results was available, through data on frequency-ratios in past trials. In such cases, action does not really involve relevant uncertainty (after all, there will be much irrelevant uncertaint in any decision); it falls within the familiar framework. Even if the particular individual will not make many trials, he can shift risk by insuring; since the relevant knowledge about many trials is accessible, institutions will be willing to specialize in combining trials; they will make a consistent profit, at rates which individuals are willing to pay. To the extent that knowledge of frequencies is available, individual may even to some extent "insure himself" by pooling all his choices in all fields.

14. But risks are "unmeasurable" when: a) no valid basis of aggregating past trials, computing frequencies; b) insuring is impossible, for that reason;(~~even~~ c) individual feels he has valid knowledge of risk, but can't communicate knowledge to an insurance institution, and will make only a limited number of relevant trials. Knight says that in this case, uncertainty, then numerical probabilities either don't exist in any meaningful sense, or they are irrelevant to problem of choice.

15. How, then, measure the likelihood of the possible outcomes? How represent the "marks" we give an outcome for likelihood? Well, on what are they based? Knight admits that a rough grouping of past instances provides a basis at least for qualitative degrees of likelihood, and describes how uncertainty can be reduced by grouping, specializing risk-bearing, etc. In short, these uncertainties seem to

much like risks; person acts pretty much "as if" he were assigning numerical probabilities to outcomes--with the qualification, which we must keep in the background of the discussion, that they are "fuzzy" at best and that they will differ from person to person. Arrow's comment (ignoring this qualification) that Knight doesn't seem to have gained much from the distinction. But Knight does not provide formal structure for dealing with individual decisions (though he describes methods for reducing uncertainty or avoiding it, institutions that are attributable to the presence of uncertainty, behavior that could not be explained without uncertainty, and the results--including profit--of the difference of behavior which would not exist without a difference in opinion and taste-for risk-bearing; note--existence of profit requires only the difference of opinion; it could exist even if one man "knew" the "correct" risks).

16. Shackle (who doesn't mention Knight) draws the line at about the same cases, describes them the same way (with a little more emphasis on the crucialness of the single decision, i.e., the irrelevance rather than the unavailability of frequency knowledge), and likewsie declares that the split corresponds to the applicability (existence and/or relevance) of numerical probability.

However, he goes further to provide a formal scheme for describing degrees of likelihood (not using numerical probs), an operation for measuring it, and a set of hypotheses for decisions based on it. Potential surprise. But where the hypothetical results of operations Knight described could be summarized by numerical probs, the results implied in Shackle's axioms definitely could not. i.e., where it would possible to describe conceivable results of Shackle's operations by numerical probs, it would not be possible with the results he predicts.

17. First, the plausibility of the notion that degrees of belief "act like" numerical probs rests on the belief that where frequency knowledge is available, the relevant degrees of belief correspond closely to these frequency-ratios (i.e., even when only a single trial is to be made). Knight clearly accepts this, in a very rough way (with great skepticism with respect to "compound" probabilities requiring sophistocated manipulation of "basic" probs). But Shackle insists that such knowledge is absolutely irrelevant to the single case, hence that degrees of belief will bear no relation to frequency-ratios. This leads him to very counterintuitive assertions about his own examples; cricket match, two decks of cards.

b) Not only are degrees of belief not based on numerical frequencies, they do not behave like numerical probs. Non-additive features. He gives cases where these results are plausible; and in such cases, we can say definitely that ~~individual does not~~ even act 2"as if" ~~he~~ assigned ~~numerical probs~~ these degrees of belief, i.e., the results of S's operation, cannot be represented by numerical probs (this does not quite settle the question whether or not the individual behaves "as if" he assigned numerical probs, ~~because~~ but in manyof the same cases, as we shall see, that is also true).

But in other cases, particular where frequency knowledge is available, these predictions about the results of applying his operation (i.e., questioning about p.s.) do not seem plausible, either. Heads or tails. Statement that relative likelihood is unaffected by rival hypotheses. (his example of land plots)

c) Shackle is probably correct that his operation would not allow us to discriminate between likelihoods of many events (coins, dice) yet other operations may make that possible, and the differences might be important. for decision-making. (again, packs of cards, coin toss)

18. Shackle really rejects numerical probs only partly for the valid reason that their formal relationships are inappropriate to the data. Partly, they reject numerical probs because of "misleading" impression of exactness; but any formalization will be guilty of that. The msin reason for their rejection is that neither Knight nor Shackle can conceive of numerical probs that are not based on frequency-ratios. They ask the sound questions: on what will these probs be based? From what information shall we (the observer) infer them? What knowledge would lead us, or the actor, to conclude that they (or our estimates) were "wrong"? What would change them?

The only answers to these questions that K and S can conceive is that numerical probs must be based on (changed by, "verified by") observation of frequency ratios, deduction from a priori probs, application of principle of indifference, etc. Hence, where this data is unavailable (Knight) or irrelevant (Shackle), numerical probs must be inappropriate.

19. In a similar attitude of despair w.r.t. numerical probs, others were looking for a way of representing uncertainty that did not utilize probabilities, and ways to describe and predict bhavior that did not imply that the individual acted "as if" he applied probs.

They found the model in the theory of games; although vN and M were motivated in constructing it by considering a different class of decisions from Knight and Shackle, and denied that it was necessary or appropriate for the cases of "games against Nature" (quote)

20. Moreover, they took over a decision-rule which did not imply that person behaved as though he assigned probs (in fact, it contradicted that); then other rules were proposed with the same nature.

Not only were degrees of belief not based on numerical frequencies, but they didn't act like probs (in fact, the notion of degrees of belief played no role in this model).

21. These models seemed appropriate, roughly, in the same cases where Shackle's rule seemed appropriate (we will see that Shackle's rule is equivalent to one of these models). But like Shackle's, there were cases where frequencies still were not available or relevant, yet where these models (like Shackle's) seemed to lose information and to fail to describe the state of mind or the decision rules. Thus, there seemed to be an area of "partial ignorance" where numerical probs might not be formally inappropriate, if they could be meaningfully defined.

22. How to measure them? a) Shackle's own operation is one method; there might well be an area within which his p.s. could be represented by probs, and these relevant to decision. More likely, there are similar operations (Coombs; "which is more likely") to provide same info. b) or like utility, since it is supposed relevant to action, measure it by means of actions (like electric current).

23. Knight himself suggested second method (quote). This idea was developed fully by Ramsey. (quote) (put up or shut up; show the courage of your convictions). Belief as a basis for action. If degrees of belief are determined in the context of action, to know what they are is the same thing as to know how they are used.

24. Ramsey-Savage method of inferring both degrees of preference and degrees of belief. In certain cases, there seems no doubt that this can be done roughly. Savage believes this this will be true for the reflective behavior of a reasonable man in all cases. If true, this wil solve problem.

25. There would remain problem of predicting non-reasonable behavior, for which other models might be appropriate. But there seems doubt that this approach will describe well even the behavior of a reasonable man in all situations. Note: where it does not apply, then person cannot even be described as acting as if he assigned numerical probs. a) can these situations be specified? b) how can you describe uncertainty and predict choice in them; will one of the game-models be appropriate?

26. Answers: a) Yes. If you insist on a purely behavioral definition, they are ~~"situations where man violates one of Savage axioms~~ "situations where it is impossible to define numerical probs such that man maxes math expectation" (it is then impossible to measure probs at all by Savage method); examples. Or, as Savage himself suggests, they are situations in which person "feels" that likelihoods are "vague," inadequate to serve as a basis of action.

27. This corresponds to an element in decision that Knight stresses very much. On the other hand, both Shackle and Savage (elsewhere) deny that person makes this two-part judgment of likelihood; they object in particular to notion that one's "confidence" in likelihoods can be expressed as a probability, since then probs could simply be compounded, and problem of confidence pushed back one stage.

But supposing that confidence is not represented in probability terms; supposing that it does not result in transforming likelihoods with a probability transformation?

28. Suppose that a person determined personal probs for himself by a Coombs-type interrogation. How then would a "lack of confidence" in these probs be reflected in his estimates: how would they be affected by doubt as to the reliability of his information or the type of grouping, by a vague expectation that new information might change them radically? I think, ~~by~~ in most cases by increasing the "likelihood" of bad outcomes and decreasing the likelihood of good outcomes; it would be these "distorted" likelihoods that would be reflected in betting behavior. Or, in cases where his attention was focussed on possible gains and the probabilities were merely unwelcome intrusions of reality, any trace of vagueness would be an excuse to exaggerate the possibilities of success (there might be considerable resistance to reducing the vagueness in this case).

Even after this has taken place, it is still possible to infer probs from a specific set of choices. But these probs would be partly dependent on payoffs. There would no longer be a set of probs independent of payoffs.

~~In an~~ extreme form, he could act as though he attached probability 1 to the worst outcome

We could frame various hypothesis as to the exact nature of this transformation; most of them would be fairly complicated and difficult to test. For practical purposes, in situations where we expect "vagueness" and hence, an influence of payoffs on probs, we could put the assumption in a radical ~~and then~~ form, which would be easy to test and apply. E.g., we could suppose that probs would depend only on payoffs; the lowest outcome might be assigned $\frac{1}{2}$, the next lowest 1/3, etc., the maximum getting the residual to sum to unity. An even simpler hypothesis would be that the minimum outcome should be assigned a probability of 1; or that the minimum should be assigned some fraction, a, and the maximum, 1-a, the person striking an individual balance between "caution" and "optimism."

29. Obviously, these decision rules will be recognized as minimax and the Hurwicz proposal, respectively. But we do not have here a general rationale for their use under uncertainty. It is suggested that they may approximate the "expectations that underlie action" in specific situations in which likelihoods are vague. And ~~even Savage~~ it is in just these situations that Savage agrees that such methods might be used (quote; actually, he endorses minimax regret).

30. We might consider this an operational definition of "vagueness"; that state of uncertainty in which we act "as if" our expectations depended on payoffs. Minimax, minimax regret, Hurwicz, other rules above are formal, approximate methods of representing this situation.

Fellner.

31. Note that vagueness is relative to payoffs. Where neither large gains nor losses are involved, we may willing to act as though our slim information were an adequate basis for precise, definite expectations. A given degree of "vagueness" will make caution (or perhaps, optimism) more compelling as the size of the payoffs increases in either direction. (Quote Savage, that his system does not allow direct role of optimism or pessimism, or confidence; this is contradicted by his discussion of vagueness; and the minimax rule, which he accepts as a response to vagueness, is simply a radical approximation to a general distortion of probs corresponding to payoffs.)

4(earlier: show how Savage system rules out probs dependent on payoffs; Suppes example))

32. Other responses to vagueness. Straight minimax and Hurwicz loses all "initial" knowledge of likelihoods. a) Hurwicz: divide outcomes into classes, eahc of which can be assigned a likelihood "confidently" but within which likelihoods are vague; use H. criterion within in each class; b) vice versa (?); c) on basis of "initial" probs, divide outcomes into "likely" and "unlikely"; treat latter as impossible, and use H. criterion ; d) Hodges and Lehmann; eliminate some as impossible, set "acceptable" minimum and max. expected utility among those which meet this minimum. e) optimistic behavior; maximax.

33. Experimental implications; probs inferred from choices involving small sums may not be reliable directly in choices involving large sums; some indications of this might show up as sums are increased (while still small); clues to the extent and direction of shift may be derived from questioning as to confidence in estimates, adequacy of information; where vagueness is suspected to be high, because of known inadequate information, suspect minimax or similar behavior. Does behavior change as "vagueness" is varied? Give subjects different information about situation (Knight example of balls in urn). i.e., increase vagueness with given payoffs. Verbal responses. Willingness to gamble. Suppose subject will take even bet on the toss of a coin; will he make the same bet if he knows the coin is biased (in an unknown fashion)? Or does he require a premium, i.e., better odds? (or will he accept worse odds?)

Or, check with utilities; compare choices with an event of"known" likelihood with one with the same but "vague" likelihood.

Or, compare results of Coombs-type experiment with probs inferred from choices; is there a closer correspondence when probs are "definite" than when they are vague relative to payoffs? If this is the case, we have means of directly estimating the effects of vagueness, comparing Coombs results with Savage results.

34. We conclude that there is a large area of uncertainty in which person may act as if he assigned numerical probs to events and numerical utilities to outcomes independently and maximized the math expectation of utility. But there is a subclss where, roughly, the uncertainty is greater--in borderline cases, this being relative to payoffs--where probs are more or less influenced by payoffs, in the extreme completely. In this latter class, a formally non-probabilistic model like minimax may serve as the most practical approximation.

35. If the distinction between risk snd uncertainty is still equated with the use of non-use of numerical probs, then the border has shifted since Knight (or Shackle). Or, we could keep dividing line between the area where knowledge of frequencies is both available and relevant, where the expectations based on them is widely and confidently acted upon--and the area where knowledge and confidence is more personal and less reliable. This last would include the actual cases that Knight classed as "uncertainty"; but we would recognize an important subdivision within that, inwhich numerical probs could not even be inferred by choices.

36. Consider games. On the one hand, we have found the game model suitable for some RobinsonCrusoe situations, where the authors believed that statistical assumptions would suffice. On the other, we have found ways to define numerical probabilities in situations where statistical frequencies were unknown or inappropriate. Perhaps vN and M ~~and~~ reject numerical probs precipitately, because they ~~associate~~ interpret them as "frequencies in long runs." Is a non-probabilistic game model really appropriate? a) Opponent's expectations are just data, though vague; ignoring that for the moment,
37. By the methods we've considered, might we assign numerical probs to opponent's strategies? Sure. Do we have definite expectations of opponent's strategies? Yes. Then consider Kaysen's proposal.

38. But: example; shift payoff to make one strategy dominant. Our expectations will depend on payoffs, because opponent's strategywill be influenced by payoffs. Examples. Our use of a mixed strategy means we are acting as if we had definite expectations of opponent's strategy. (Remember, these particular "as if" hypotheses didn't come out of the blue; they are based on introspective evidence as to actual mental processes). Familiar method of measuring utility by considering how often person chooses act; notice the paradoxical results; increasing a payoff may make the act less popular (contrary to Kaysen); may shift us from one act to another. The explanation is that changing the payoffs automatically shifts our expectations.

Likewise, adding or subtracting an irrelevant act may change our choice.

If there is a reason to suppose opponent will choose ~~xxx~~ strat A... and such supposes will probably start with a payoff; opponent will be aware that we will start there; then there will be reason for me to use B; does consideration of this make it more likely that he will use A? If it does, then there is double ground for the expectation. If it doesn't (it won't unless both are minimax) (rather, close to minimax) will this expectation have a clear effect on opponent's choice (mutually recognized)... Process needn't go on till "expectations are coordinated"; in practice, players will make differing judgments of opponent's expectations.(which may be definite and confident)

39. Conclusions: the essential difference between games against Nature and games against an opponent is that expectations will definitely depend on payoffs. Moreover, we can predict the nature of the relation-

plausibility considerations are less important when direct testing is possible; hence we need be less concerned over whether a hypothetical use of minimax violates certain postulates

ship: a) the difficulty of estimating the opponent's expectations will often result in vagueness, which will usually result in "conservative" distortion of likelihoods (note: it might have opposite effect).
b) Knowledge that opponent will likewise suffer from vagueness makes conservative behavior on his part more likely, which will reinforce ~~xxxxxxxxx~~ one's own minimax behavior--with reciprocal effect.
c) A low payoff for one strategy means a high one for the opponent; we will exoect that the corresponding strategy will at least draw the favorable attention of the opponent--making it "potentially more likely." Will opponent think of this thought process on our part? Will the potential losses of that strategy scare him off?
d) the possibility of being found out will make bad outcomes more likely; this may sometimes be a definite expectation; and the

resulting expectations are mutually reinforcing.

40. Thus, introspective considerations lead us to reject the method of measuring expectations in a game by varying the payoffs. This leaves the possibility of measuring them by direct methods for particular games, checking whether the influence of payoffs is in the direction indicated. This can lead to hypotheses directly on game behavior. One of these might well be the game solution: justified here as an approximation to a model in which expectations are influenced in a general way by payoffs.

41. In non-zero sum games, expectations are in the center; Luce example of "psychological dominance."

42. The attitude here is that rational conduct in the game situation, as in any other, depends on expectations; though it happens that a formally non-probabilistic model may be a practical solution of representing the interrelationships of expectations and payoffs. If definite expectations are non-vague, and if they survive consideration of the possibility that that the opponent may foresee them with high confidence, then I think that players will and indeed should act on them. To go further: while it may be true in a ~~xxxx~~ zero-sum game that the better certain payoffs look, the less likely they willbe achieved--this does not follow with the mechanical rigor implied by the strict minimax solution; in some circumstances, where both sides see the possibility of payoffs much better than minimax, then anything less than certainty that the opponent will minimax would justify abandoning minimax. (my example) (thus, this argument against minimax is not as strong as inthe game against Nature, but it still has some force).

43. Finally; ~~xpxx~~ we said problem was one of organization choice. This has the nature of group choice, where members of the group have same preferences (we will assume, even for risks). But if situation involves uncertainty, then the first indication of this will be differing estimates of outcomes. General problem of group choice; quote Stigler, Friedman.

a) individual or group uncertainty can result in agreement on conservative estimate.

b) Savage; minimax regret sol (with possibility of experiment).

expect. clearly enter in defining "possible" strats plausibility considerations – axioms

Sure-thing Principle, Rubin Ax.

Relation of Shackle – Hurwicz

Emphasis on empirical implications of hypotheses.

Comments on empirical measurement of prob, utility. Ordered metric.

Axioms on Total Ignorance

Stock market